Strategic approach to IT scalability

By Pawel Plaszczak, GridwiseTech

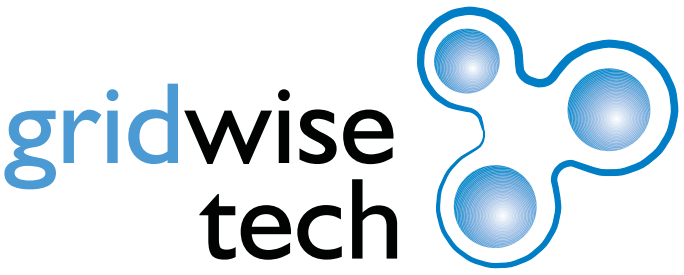

GridwiseTech
ul. Chrobrego 28/4, 31-428 Kraków, Poland
Phone: +48 12 294 71 20

www.gridwisetech.com
contact@gridwisetech.com

1. What is scalability?

A bank closes for customers four times a year, due to periodic processing.

An entertainment portal makes customers wait for minutes for the movie download to complete. In the meantime, the competition streams the data in real time.

An automotive manufacturer postpones the new car model launch date because of delays in auto body rupture prediction simulations.

Is there a common pattern? Naturally, all three situations are unwelcome. Also, each of them could be avoided if handled by a scalable IT architecture.

What is scalability? In the traditional sense, a scalable system increases efficiency in proportion to the resources available. If a company uses a multi-processor server rack to serve an average customer in 30 seconds, purchasing another rack should bring the response time down to 15 seconds. Although such behavior seems intuitive, this is by no means inherent in most of today's IT systems.

There is also a modern notion of „dynamic" scalability: a scalable system will always maintain the requested quality of service. Thus, if the demand exceeds the capabilities of the available hardware, a scalable system will dynamically allocate extra resources in (or even outside) the company, to process the workload in the requested time. Imagine a company with two servers. The first server hosts a priority service for customers with a 15-second guaranteed response time. The second server hosts back office applications, with no constant availability requirements. When the first server faces a momentary peak in customer demand, the scalable system will allocate extra processors from the second server to help with priority processing. The service-level agreement (SLA) is maintained, customers are satisfied and no unnecessary hardware expenses have been made.

Both „traditional" and „dynamic" scalability serve similar purposes. Scalable systems are designed primarily to minimize IT investments and increase market readiness.

2. Scalability problems will become more severe

Many of today's production systems (outside large data centers) do not even have traditional scalability properties, nor do they embody the dynamic scalability principle. It is interesting to see why such a pattern prevails, at a time when scalable technology is not only needed, but also widely available. The reason is that we are in a transitional period. Even ten years ago, scalable systems were not always necessary. Today the situation has changed dramatically, while businesses still attempt to struggle with their old architectures. A short look at the market trends reveals that things will only get worse for them.

In the last few decades, the market has discovered virtual objects, cheaper to operate than the physical ones. An electronic book costs almost nothing to maintain, transport, search through, or copy. A virtual car model is less expensive to crash than the real one. Virtual protein docking is not only faster but also raises fewer concerns than testing of potential drugs on animals. And so on. If virtualization was a sporadic phenomenon in the 1990s, it is a massive trend today. This means more data, processing and traffic, and in general – an increase of scalability demand. And the pace is scary. In 2000, Wachovia was the only bank to interconnect the first 250 desktops in a department grid. In 2007, the size of their data centers is essential for derivative market competitors, with leaders such as Citigroup and Fortis each running 10,000-CPU grid.

Secondly, the global digital market continues to grow. In China, 2006 brought a 23% increase in the number of Internet users. Ten years ago the money spent on www during the electorial campaign could be money lost. Today, parties lose election battles due to a low online marketing budget. Furthermore, radio stations hastily change their mode of operation to become interactive, presenting online news, music and media shops. They have a good reason to fight – in the saturated, mature market of radio broadcasters, every percent of the market must be wrested from the competition in a fierce battle. At the same time, a well-marketed online broadcasting service can triple the number of listeners in a few weeks. As a consequence of the ballooning market, the network traffic of online businesses grows inexorably.

There are also more complex reasons to believe that scalability needs will continue to grow. To name only one, we are in the epoch of global-scale mergers and acquisitions. Mergers often entail the need for higher throughput systems. On top of all that, one should remember the add-on data phenomenon: data that was self-sufficient in paper form often generates additional data (read: network traffic) when transferred to digital form.

The above means grave consequences for business. If your system faces scalability problems today, it will only get worse tomorrow. What is more, we are in a self-perpetuating spiral of traffic, data and processing that generates more traffic, data and processing. The trend will continue in the coming decade, while the more distant future is still beyond imagination.

Conclusion: the question of IT system scalability is inevitable for any business in search of a successful and sustainable market presence.

3. The strategic approach

Can the scalability problem be approached in an organized way? Is it possible to foresee scalability issues before they occur, and provide a remedy beforehand? Our vision is limited to the age we live in, and thus, today's best answer may not be optimal in ten years. However, thanks to the pioneering work that has already been done in this area, some generic guidelines can be given.

Having analyzed and understood the market and technology trends that give rise to scalability issues, we have identified the bottlenecks. Before focusing on IT, it is important to understand that IT aspects may only be a fraction of the problem. Operational scalability may also be limited by a spectrum of legal, administrative and procedural obstacles, lack of flexibility in business processes, and the human factor. From the technical standpoint, the scaling of IT must proceed in line with the scaling of logistics, infrastructure, and staff management.

When it finally comes to IT issues, one must consider scalability at each layer of technology. Building

detailed event scenarios helps identify the layers in question. The weakest link principle applies here: a bottleneck hidden in one technology layer will foil endeavors to increase the throughput and efficiency of other layers. To be more precise, in a typical system one must look separately at:

Application-specific logic that must be designed with high throughput use in mind. Monolithic applications often need to be broken into parts, with the portal-based presentation layer physically detached from the processing-intensive business logic routines.

Web layer and application servers to handle the fluctuating irregular load of external requests and transparently pass them further

Service layer (SOA) that must be able to dynamically adapt the loosely coupled architecture to a brief temporary demand increase

Execution layer, including distributed Grid and cluster parallel processing. Grids can greatly relieve the workload from the Web layer and manage load balancing over all available resources

Data access and database layer, where bottlenecks can be avoided through intelligent caching and distribution

One should not forget the vertical cross-cutting technology concepts: security, network architecture, management, adaptability and virtualization paradigms.

4. In practice

The ideal outcome of the scalability work is a fully virtualized infrastructure, with resources dynamically available in an on-demand model. In most cases, however, even partial virtualization is enough to satisfy performance requirements. The concept can be realized by detailed work focused in the layers identified as critical. We can explain this in detail by using examples from the author's business practice.

An engineering sector manufacturer needed to perform advanced simulations. Their product had to go through complete computer simulation testing before the production of the physical prototype. Each simulation cycle took thousands of CPU-hours and produced terabytes of intermediate data. The company was looking to improve their time to market (TTM) by shortening the design cycle by months. As the lacking resource was the computation time, the solution involved scaling out the workload into the grid layer. In the improved system, simulations are managed by the grid execution management module, which automatically allocates jobs to CPUs. External computing resources were used to serve peak demand. Distributing the application also involved application-scope and Web layer work. The simulation interface is now available to users through the browser. By just one mouse-click, the task is transparently distributed to dozens or hundreds of servers in a way invisible to the user. After a few hours (instead of months), the user can inspect the results by means of the same interface.

A B2B company dealing with data exchange among retail businesses expects to process one hundred million documents a year, containing sensitive data of individuals. The response time must be in seconds. Are we facing a scalability challenge?

A quick calculation shows that the company needs to process ten documents a second (assuming the working hours only). As a typical email server is able to send up to 50,000 emails per hour, the number does not look impossible. A further analysis of the time distribution of workload should tell whether the processing needs improvement.

However, looking at the global picture, one will notice a more severe scalability issue in the domain of security. How to ensure universal document protection when communicating with a large number

of external institutions? One must secure trusted communication among partners, which is typically achieved by means of certificates. However, the certificate-based approach encounters the adoption barrier, especially if each employee of all the communicating parties is required to possess a certificate. Thus, the approach does not scale. A federated authentication paradigm could be deployed to ensure identity protection.

In the banking sector, a frequent issue is the end-of-quarter processing of all accounts, often running hours behind schedule and causing operational hindrance. In this case, bottlenecks are usually close to the database. It is not the computing itself, but rather the database management system that is causing the end-of-quarter process to run out of control.

Most banks deal with the issue by purchasing new rows of expensive hardware to deal with the problem. However, this can often be replaced by an inspection of the suboptimal database application code that was written before the scalability issue arose. As an alternative, both banks and vendors look at data caching solutions that enable the distribution of the data access layer. Thus, the bottleneck is effectively exported from the database domain into the distributed processing domain, where it can be dealt with parallel processing methodology.

The three mentioned examples show how different types of business differ in scalability requirements. As requirements for scalability will grow in most business domains, forward-thinking enterprises are concerned with IT scalability in all technology layers, spotting the issues in time before the business efficiency is affected. The complex approach helps, and an analysis of the event scenarios helps identify priorities in the scalability improvement process.

Pawel Plaszczak's

international software engineering experience includes CERN, British Telecommunications and Argonne National Laboratory. Having earned his Msc in Computer Science in 2000 from AGH University of Science and Technology, he later worked for Ian Foster's Globus Project, fundamental in the history of Grid computing movement. In 2003 Pawel Plaszczak left Globus to found GridwiseTech and work for customers who are the earliest adopters of Grid technologies. He is the author of numerous articles and tutorials, as well as the book "Grid Computing: The Savvy Manager's Guide", explaining Grid computing in business terms.